

Privacy-Preserving Public Information for Sequential Games

AVRIM BLUM, Carnegie Mellon University
 JAMIE MORGENSTERN, Carnegie Mellon University
 ANKIT SHARMA, Carnegie Mellon University
 ADAM SMITH, Penn State University

In settings with incomplete information, players can find it difficult to coordinate to find states with good social welfare. For example, in financial settings, if a collection of financial firms have limited information about each other's strategies, some large number of them may choose the same high-risk investment in hopes of high returns. While this might be acceptable in some cases, the economy can be hurt badly if many firms make investments in the same risky market segment and it fails. One reason why many firms might end up choosing the same segment is that they do not have information about other firms' investments (imperfect information may lead to 'bad' game states). Directly reporting all players' investments, however, raises confidentiality concerns for both individuals and institutions.

In this paper, we explore whether information about the game-state can be publicly announced in a manner that maintains the privacy of the actions of the players, and still suffices to deter players from reaching bad game-states. We show that in many games of interest, it is possible for players to avoid these bad states with the help of *privacy-preserving, publicly-announced information*. We model behavior of players in this imperfect information setting in two ways – greedy and undominated strategic behaviours, and we prove guarantees on social welfare that certain kinds of privacy-preserving information can help attain. Furthermore, we design a counter with improved privacy guarantees under continual observation.

1. INTRODUCTION

It is widely accepted that one cause of financial crises is many players (such as investment banks) making risky and highly correlated investments. In some cases, this concentration of risk can be intentional—all players believe that these actions will be profitable, even knowing how crowded the market has become—but in other cases this clustering appears to be due to a *lack of information*¹. Researchers at federal agencies are actively investigating methods to publish useful information about the state of markets in the hope of diffusing future crises (e.g. Oet et al. [2012]). Even though regulators have access to detailed confidential information about financial institutions and (indirectly) individuals, current statistics and indices are based only on public data, since disclosures based on confidential information are restricted. However, forecasts based on confidential data can be much more accurate², prompting regulators to ask whether aggregate statistics can be economically useful while also providing rigorous privacy guarantees [Flood et al. 2013]. Suppose that a clearinghouse agency indeed had information about the different actions made by players so far: *could it, in a way that preserves privacy of the actors, release enough information to help avoid disastrous outcomes?*

In this paper, we consider an online, multi-agent decision-making setting in which we can pose and formally analyze this question. In this setting, a clearinghouse agency may publish *differentially private* information about actions of agents so far, with a goal of *improving social welfare* or, at the very least, *avoiding disastrous sequences of decisions*. While we use the financial setting to motivate our research question, the

¹For example, the Financial Crisis Inquiry Commission [2011, p. 352] concludes that, “The OTC derivatives market's lack of transparency and of effective price discovery exacerbated the collateral disputes of AIG and Goldman Sachs and similar disputes between other derivatives counterparties.”

²For example, Oet et al. [2012] compared an index based on both public and confidential data with an analogous index based only on publicly available data. The former index would have been a significantly more accurate predictor of financial stress during the recent financial crisis (see Oet et al. [2011, Figure 4]). See Flood et al. [2013] for further discussion.

question itself is relevant in many contexts as illustrated by some of the related work (discussed further in Section 3).

1.1. An illustrative setting

Consider the following abstract financial decision-making model. There is a set of m “markets” or “resources” (conceptually, view these as nodes on the right-hand-side of a bipartite graph) and n players (view these as nodes on the left-hand-side of a bipartite graph) who will arrive online, one at a time. Each player i has some set A_i of allowable actions to choose from, known only to that player. In the simplest case, A_i will be just a set of markets (equivalently, a set of edges incident to player i in the bipartite graph) and the action of player i will be just to choose one market in A_i ; more generally, A_i will be a set of “investment portfolios” available to player i , where each investment portfolio is itself some (potentially fractional) allocation among the markets.

The players make their investment choices in some *arbitrary sequential order*. Suppose each market r has some non-increasing function $v_r : \mathbb{Z}^+ \rightarrow \mathbb{R}^+$ indicating the *value, or utility, of this market to the k th player who chooses it*. For example, one might have $v_r(k) = v_r^{\text{init}}/k^p$, for $p > 0$, where v_r^{init} is the initial value of market r , so that the value of market r to new players rapidly drops as a function of the number of players who have chosen it so far³.

Let us consider the game-play under two extreme information settings.

- (1) *Perfect information*: In this setting, if each player i has *perfect information* about the investment choices made by the players before her, the optimal action for player i is to greedily select the action in A_i of highest utility based on the number of players who have selected each market so far. Moreover, it is not hard to show that if players behave in this manner, this will result in *total social welfare within a factor of 2 of social optimum*.⁴
- (2) *No information*: Suppose, however, that the player i has *no information* about the investment choices made by earlier players. In addition, she does not know what markets other players are interested in and even how many players have played so far. The only information she has is her set A_i and knowledge of the functions v_r . In this *no information* case, some particularly disastrous sequences of actions might reasonably occur, leading to very low social welfare. For example, suppose each A_i contains markets r and r' , with $v_r(0) = 1$, $v_r(k) = 0$ for all $k \geq 1$, but $v_{r'}(k) = 1/2$ for all k . Without additional information, players might reasonably choose greedily according to $v_r(0), v_{r'}(0)$, selecting the market of higher initial value. This would give social welfare of 1, whereas the optimal assignment would give $n/2 + 1/2$.⁵ Even worse, suppose each A_i contains one market r_i with $v_{r_i}^{\text{init}} = n - 1$ and r with $v_r^{\text{init}} = 1$, and each market behaves as $v_q(k) = v_q^{\text{init}}/k$ for all k, q . If each player believes (incorrectly) that the high-value market is shared and has already been chosen by all the $n - 1$ other players, and so chooses the low-value market. In this case, if in fact it is only the low-value market that is shared, *the total social welfare will be approximately $\ln n$ when it could have been $\Omega(n^2)$* .

As the above example illustrates, having information about the current state of the system can be important to achieving a good social welfare. Suppose that a clearing-house agency were able to collect such information about the actions taken by players

³In Section 6, we consider the setting in which payoff is affected also by actions of *future* players.

⁴Please see Section 5 for details.

⁵One can think of this scenario as a case where each player i has a bank account r' and is tempted by an investment r that offers the potential to double his money. However, if all players choose it, only the first one actually doubles his money, and the rest lose all their net worth.

so far, but it was required to satisfy strict privacy conditions on any data it published. Could it avoid the above kinds of disastrous event sequences? More specifically, the question we consider is:

If a clearinghouse agency can only post differentially-private⁶ information about the actions of players, can this suffice to provide strong guarantees on social welfare?

What we show is that in a quite general sense, the answer to the above question is yes. In fact, in some cases, providing private information can achieve social welfare that exceeds even the full information setting (see Section 8.3). The central agency need only observe actions taken by players over time and post differentially private estimates of the current usage of each resource. Furthermore, we design new privacy-preserving counter mechanisms that are particularly well-suited for this task, combining additive and multiplicative guarantees in order to be especially accurate when we need accuracy most.

To help analyze game dynamics, we consider two models on player behavior:

- (1) **Greedy behavior:** here we assume players act greedily, selecting the highest-utility action according to the estimates (recall that this would be *optimal* behavior if the estimates were perfect, so this is natural behavior to consider), or
- (2) **Behavior in undominated strategies:** here we make only the minimal assumption that players will not choose dominated actions (if action a is guaranteed to be worse than action b in all states that are possible given the information provided, then action a will not be chosen) and consider the worst-case sequence of actions subject to this mild constraint.

We prove strong guarantees on social welfare under differentially-private information for both models (stronger, of course, for greedy behavior, which is a specific undominated strategy). Note that positive worst-case guarantees for undominated strategies (worst-case over the sequences of sets A_i and over the undominated strategies chosen) imply that even players with wildly inaccurate beliefs will produce good social welfare. One could also imagine less pessimistic models such as Bayes Nash equilibria. We discuss these other possibilities in Section 9. Furthermore, as discussed in more detail in the same section, we note that, while the privacy-preserving information provided to the players will give the approximate usage of each resource and this the players shall use to choose their action, it *not necessarily* the case that the value received by each player will be close to what that she could received had she been given exact usage estimates, i.e., *approximate usage estimates do not imply approximately optimal investment decisions*.

2. STATEMENT OF MAIN RESULTS

For the sequential resource-sharing games that was described above, we prove that with privacy-preserving information and under the greedy strategy, the competitive ratio⁷ is bounded and is *polylogarithmic* in the number of players. These results should be contrasted with the results on perfect information and no information cases discussed briefly in the introduction and in more detail in Section 5. For undominated strategies, we also prove a polylogarithmic guarantee for resource-sharing games whose values drop at a polynomial rate ($v(k) = v^{init}/k^p$ rate for constant $p > 0$), but show that if values may drop arbitrarily fast then the competitive ratio is unbounded. Resource sharing games have the property that the utility of a player depends on her

⁶Differential privacy [Dwork et al. 2006] is a well-studied and strong formal privacy notion, see Section 4.

⁷We use the usual definition of competitive ratio; for a formal definition, see Section 4

action and *only* the actions taken by *previous players* in the sequence. To understand how crucial this is, we consider market sharing games and a generalization of them where players' utilities also depends on *actions taken by future players* (Section 6). While market sharing games allow for a logarithmic competitive ratio with greedy strategies, their generalization do not have allow for a bounded competitive ratio even with greedy strategies. Certain classes of valuation functions that decrease gradually have a bounded competitive ratio.

The key privacy tool we use is the differentially private counter under continual observation [Dwork et al. 2010]. Incorporating these counters into the sequential games maintains the privacy of the players' information, and we show that compared to the case of having exact (non-private) counters, the degradation in the competitive ratio for the resource sharing game is polylogarithmic in the number of players. In addition, we also improve upon the existing error guarantees of differentially private counters and design a new differentially private counter in Section 7. The new counter provides a tighter additive guarantee at the price of introducing a constant multiplicative error.

In Sections 8.1, 8.2, and 8.3, we consider other classes of games – specifically, we analyze Unrelated Machine Scheduling, Cut, and Cost Sharing games. The work of Leme et al. [2012] showed these games have improved *sequential* price of anarchy over the *simultaneous* price of anarchy. For these games, we ask the question: if players do not have perfect information to make decisions, but instead have only noisy approximations (due to privacy considerations), does sequentiality still improve the quality of play? We prove that the answer is affirmative in most cases, and furthermore, for some instances, having *differentially-private information dissemination improves the competitive ratio over perfect information* (Proposition 8.8).

Of independent interest, we show an extension of the well-known 2-approximation which greedy achieves for online vertex-weighted matching. We show that the greedy algorithm 4-approximates OPT for a general class of many-to-one vertex-weighted matching settings. Suppose nodes on the left arrive online, with *allowable subsets* of nodes on the right-hand side of the graph. If these sets are downwards closed (e.g., that a player allowed to take a set S implies they are allowed to take any $S' \subseteq S$), then Greedy is a 4-approximation to OPT . See Section A for details.

3. RELATED WORK

A great deal of work has been done at the intersection of mechanism design and privacy; Pai and Roth [Pai and Roth 2013] have an extensive survey. While the survey exhaustively covers the research work at the intersection of the two areas, we list here a few key connections. One line of work has employed techniques from differential privacy to design incentive compatible mechanisms (for e.g., [McSherry and Talwar 2007; Nissim et al. 2012b]). More recently, differential privacy has been considered as a constraint, and incentive compatible mechanisms that satisfy this constraint while optimizing an objective function have been designed [Huang and Kannan 2012; Xiao 2013]. In other work, the utility function of agents have been redesigned to incorporate privacy concerns, and mechanisms for the redesigned utility functions have been constructed [Chen et al. 2013; Ghosh and Roth 2011; Roth and Schoenebeck 2012; Ligett and Roth 2012; Fleischer and Lyu 2012; Nissim et al. 2012a].

Our work is similar to some of the previous work in that it considers maintaining differential privacy to be a constraint. The focus of our work however is on *how useful information can be provided to players in games of imperfect information* to help achieve a good social objective while respecting the privacy constraint of the players. The work of Kearns et al. [2012] is close in spirit to ours. Kearns et al. [2012] consider a game where players have incomplete information, and provide a mechanism that helps implement an equilibrium by collecting the information from the players and

providing them suggestions on what actions to take. The mechanism is designed so that it is incentive compatible for the players to participate in the mechanism, and to follow its suggestions. Other related work includes that of Rogers and Roth [2013], which shows how to privately compute approximate Nash Equilibria in congestion games. In the context of interactive mechanisms, other work [Hsu et al. 2013] shows how to privately compute approximate Walrasian equilibria.

Another piece of work that is relevant to us is that of Leme et al. [2012]. While the paper focuses on bringing out the differences between *sequential* and *simultaneous* versions of certain games, their work can also be reinterpreted to be asking how providing *complete information* about the state of the game can help players achieve a good social objective in sequential games. The work examines the quality of the Nash equilibrium when players play sequentially and contrasts it to the case when the same game is played simultaneously. The main message of the work is that in certain games, playing *sequentially with complete information* helps the social objective as the players playing later keep in mind the moves made by the earlier players. In our work, we consider some of these games; however, our results highlight the games in which the information provided to the players is made differentially private either allows for improvement over simultaneous PoA or does not.

As mentioned in the introduction, one class of player behavior under which we analyze the games is *greedy*. Our analysis of greedy behavior is in part inspired by the work of Balcan et al. [2009], who study best response dynamics with respect to noisy cost functions for potential games. They ask: if players begin at some initial state, how much can the social welfare degrade after a sequence of best-response moves made with respect to noisy or slightly perturbed costs? A notable distinction between their setting and ours, however, is that the noisy estimates we consider (and that differentially-private counters provide) are *estimates of state, not value*, and may for natural value curves be quite far from correct in terms of the *values* of the actions⁸.

In a later section of this paper, we consider the case where players' utilities depend on the actions of all other players. These games are a generalization of market-sharing games⁹. Previous work [Goemans et al. 2004] shows the Price of Anarchy of market-sharing is 2, that players selecting β -approximately greedily in a single round of best-response dynamics achieve a competitive ratio of $\Theta((\beta + 1) \log(n))$. Market-sharing games are a special case of congestion games, whose price of anarchy has been studied extensively [Christodoulou and Koutsoupias 2005a,b; Roughgarden and Tardos 2004; Roughgarden 2003; Suri et al. 2007], and where greedy strategies have been shown to have constant competitive ratio in special cases [Suri et al. 2007; Awerbuch et al. 1995].

4. PRELIMINARIES

4.1. Game Model

Consider the setting in which there are m resources and n players. An action a_i of player i is of the form $(a_{i,1}, \dots, a_{i,m})$, where $a_{i,r}$ represents the amount that player i 'invests' in resource r . We denote by A_i the set of all possible actions a_i that player i can take. The players will arrive one at a time, in an order $\sigma : [n] \rightarrow [n]$ which is *adversarially chosen*. Suppose, for ease of exposition, we rename players such that player i is the i th to arrive. For simplicity, we first assume that all $a_{i,r} \in \{0, 1\}$ (for the continuous version, see Section D).

⁸Having a 'good' estimate for x does not necessarily imply a 'good' estimate of $v(x)$ for many natural value curves $v(\cdot)$.

⁹In the case of market-sharing, a market r has a fixed value v_r , which is shared equally amongst players who service the market; we present a generalization of these games. Players can fractionally service markets, and for the restrictions on collections of markets can be arbitrary rather than budgetary.

We consider two classes of utility functions for the players.

- (1) **Resource Sharing:** In this setting, the utility to a player of choosing a certain resource is a function of the resource and (importantly) only of the number of players who have invested in the resource before her. In this case, each resource r has some *non-increasing* function $v_r : \mathbb{Z}^+ \rightarrow \mathbb{R}^+$ with $v_r(k)$ indicating the value, or utility, of this resource to the k th player who chooses it (more generally, if we allow choosing resources with fractional allocations, we will have $v_r : \mathbb{R}^+ \rightarrow \mathbb{R}^+$). Let $x_{i,r} = \sum_{j=1}^{i-1} a_{j,r}$ for each r . Then, the utility of player i is $u_i(a_i, a_1, \dots, a_{i-1}) = \sum_r a_{i,r} v_r(x_{i,r})$.
- (2) **Future dependent setting:** Here the utility to a player of investing in a particular resource is a function of the total number of players who have chosen that resource, including those who have invested after her. In this case, if k players have invested in the resource, each of them receive $v_r(k)$ utility. The utility of player i is then $u_i(a_1, \dots, a_n) = \sum_r a_{i,r} v_r(x_{n,r})$ with $x_{n,r} = \sum_{j=1}^n a_{j,r}$ for each r .

4.2. Information Model

Privacy-preserving public announcements: We will be interested in designing announcement mechanisms \mathcal{M}_i which can intuitively be thought of as giving some information about actions made by the previous players to player i . Furthermore, we will not assume that players have any other knowledge about either the game play or the types and the strategies of the other players.¹⁰

Mechanism $\mathcal{M}_i : ([0, 1]^m)^{i-1} \times R \rightarrow H$ (for H some output space) can depend upon the actions taken by players until timestep i (but not the types of any players), and potentially on internal random coins R . When player i arrives, $m_i(a_1, \dots, a_{i-1}) \sim \mathcal{M}_i(a_1, \dots, a_{i-1})$ will be publicly announced. Player i plays according to some strategy $d_i : H \rightarrow A_i$ unknown to the announcement mechanism. She applies d_i to the announcement when it is her turn to choose her action $a_i = d_i(m_1, \dots, m_i(a_1, \dots, a_{i-1}))$, a random variable that is a function of this announcement. When it is clear from context, we will denote $m_i(a_1, \dots, a_{i-1})$ by m_i . We now define an (ϵ, δ) -differentially private announcement mechanism.

Definition 4.1. An announcement mechanism \mathcal{M} is (ϵ, δ) -differentially private under adaptive¹¹ continual observation in the strategies of the players if, for each player i , each pair of strategies d, d'_i , and every subset S of events in the output announcement space H^n :

$$\mathbb{P}[(m_1, \dots, m_n) \in S] \leq e^\epsilon \mathbb{P}[(m_1, \dots, m_i, m'_{i+1}, \dots, m'_n) \in S] + \delta$$

where $m_j \sim \mathcal{M}_j(a_1, \dots, a_{j-1})$ and $m'_j \sim \mathcal{M}_j(a_1, \dots, a_{i-1}, a'_i, a'_{i+1}, \dots, a'_{j-1})$. Here $a_j = d_j(m_1, \dots, m_j)$, and $a'_i = d'_i(m_1, \dots, m_i)$, and for all $j > i$, $a'_j = d_j(m_1, \dots, m_{i-1}, m_i, m'_{i+1}, \dots, m'_j)$.

In other words, we require that in the two worlds differing in a single player changing her strategy from d_i to d'_i , the joint distribution all the players' announcements (and thus their joint distributions over actions) should be statistically close. Note that the distribution of each person's announcement after i may change slightly, causing their actions to change slightly; this necessitates the cascaded m'_j, a'_j for $j > i$ in our definition.

¹⁰If players have any additional information, they are free to use it while making their choice in the case of undominated strategies.

¹¹Adaptivity is needed in this case because the announcements are arguments to the actions of players: when a particular action changes, this modifies the distribution over the future announcements, which in turn changes the distribution over future selected actions.

Privacy-preserving public announcement mechanism used in the paper: Our announcement mechanism will maintain a differentially private counter for each resource r . We will refer to this set of counters as *counter vector*. The counter vector will (approximately) track the amount of use that the resources have received over time. The values of all the counters in the counter vector are *publicly announced* throughout the game play. They will guarantee privacy of the players under adaptive continual observation as defined in Definition 4.1. Each player observes the current values of these counters when she makes her decision to invest in the various resources. Since a counter vector will need to maintain differential privacy while giving out close estimates of the usage, we now define what we mean by accuracy of a counter vector.

Definition 4.2. The counter vector will be defined to be (α, β, γ) -accurate if with probability at least $1 - \gamma$, at all points of time, the displayed value of all the counters y_i lies in the range $[\frac{x_i}{\alpha} - \beta, \alpha x_i + \beta]$ where x_i is the *true* count for resource i ¹².

We will refer to $(\alpha, \beta, 0)$ -accurate counter vector as (α, β) -counter vector for brevity.¹³ (It is possible to achieve $\gamma = 0$ ¹⁴, taking an appropriate loss in the privacy guarantees for the counter (Proposition 7.4).)

To help contrast the dynamics of the game with differentially private counters, we will additionally analyze the game under two extreme settings of counters. With *perfect counters*, the counter for each resource, at all points in time, displays the *exact count* of the number of players who have chosen the resource. With *empty counters*, the players are provided no information regarding how many players have chosen a particular resource. For ease of exposition, we assume empty counters are identically 0 always (and for the analysis of greedy behavior, this is equivalent to assuming players choose according to $v_r(0)$).

To summarize, we will consider 3 sorts of announcements: perfect counters, empty counters, and (ϵ, δ) -privacy preserving (α, β, γ) -accurate counter vector. By perfect counters, we mean, $m_i(a_1, \dots, a_{i-1}) = (x_{i,1}, \dots, x_{i,m})$; empty counters means that $m_i(a_1, \dots, a_{i-1}) = \vec{0}$; and differentially private counter vector $m_i(a_1, \dots, a_{i-1}) = (s_{i,1}, \dots, s_{i,m}) \in [0, n]^m$ where $s_{i,r}$'s are approximations to $x_{i,r}$'s. An implementation of an (α, β, γ) -accurate counter vector that is (ϵ, δ) -differentially private is discussed in Section 7.

4.3. Players' Behavior

Since we consider game play with incomplete information, we need to describe the strategy space of the players. We will analyze the game play under two classes of strategies – greedy and undominated strategies. We assume that the value functions v_r for all r are known to all players.

- (1) **Greedy strategy:** Under the greedy strategy, a player chooses the resource that maximizes her utility given the *currently displayed (or announced) values of the counters*¹⁵. This means that when player i gets her turn to play and say, $m_i(a_1, \dots, a_i)$ is $(s_{i,1}, \dots, s_{m,i})$, she picks the resource r with largest $v_r(s_{i,r})$. Such a player has no outside information or belief about the game or the types and the strategies of the other players.

¹²The counter we define has this hold for a collection of m counters simultaneously; with probability $1 - \gamma$, all counters will be within their appropriate range.

¹³We note that, for our results regarding greedy play, a tradeoff between ϵ, δ , and γ occurs.

¹⁴This is necessary for undominated strategies, which will assume the multiplicative and additive bounds on y are worst-case.

¹⁵Even with a *future-dependent* utility function, the player will choose the resource that will maximize her utility function given the *current counter* values during her turn (and not concern herself with what future players might choose).

- (2) **Undominated Strategy (UD):** A player is allowed to play *any* undominated strategy a_i with respect to the announcements m_i . A strategy a_i is *dominated* with respect to m_i if there is some a'_i for which $u_i(a'_i, a_{-i}) > u_i(a_i, a_{-i})$ for all a_{-i} which could cause m_i with positive probability. If no such a'_i exists, a_i is said to be *undominated* with respect to m_i . For instance, in the resource sharing game, with $(\alpha, \beta, 0)$ -accurate counter vector, this implies that picking resource r for a player is dominated if there exists resource r' such that $s_{i,r'}/\alpha - \beta > \alpha s_{i,r} + \beta$, (for $v_r[\cdot] = v_{r'}[\cdot]$), under undominated strategic play, a player can pick any undominated resource.

4.4. Competitive ratio

We analyze the social welfare $SW(a) = \sum_i u_i(a)$ generated by an announcement mechanism \mathcal{M} and compare it to the optimal social welfare OPT . For a game setting g , constituted of a collection of players $[n]$ and their allowable actions A_i (as defined in Section 4.1), $OPT(g)$ is defined as the optimal social welfare that can be achieved by any allocation of resources to the players, where the space of feasible allocations is determined by the setting g .

Our focus in this paper will be to investigate the social welfare guarantees achieved with perfect, empty, and differentially-private counter vector under greedy or undominated strategic play. We will denote by $CR_{\text{GREEDY}}(g, \mathcal{M})$ and $CR_{\text{US}}(g, \mathcal{M})$ to denote the ratio of $OPT(g)$ to the social welfare achieved under greedy and undominated strategies respectively, using mechanism \mathcal{M} . In case \mathcal{M} uses internal random coins, our results will either be worst-case over all possible throws of the random coins, or will indicate the probability with which the social welfare guarantee holds. We consider perfect counters ($\mathcal{M}_{\text{Full}}$), empty counters (\mathcal{M}_{\emptyset}), and differentially-private counters.

5. RESOURCE SHARING

In this section, we consider resource sharing games – the utility to a player is completely determined by the resource she chooses and the number of players who have chosen that resource before her. This section considers the case where players' actions are discrete: $a_i \in \{0, 1\}^m$ for all i , $a_i \in A_i$. We defer the analysis of the case where players' actions are continuous to Appendix D.

We first present results for the cases of perfect counters and empty counters. For the case of perfect counters, we prove that playing the greedy strategy strictly dominates any other choice of action, and therefore, under undominated strategic play, each player shall still play only the greedy strategy. For perfect counters, therefore, greedy and undominated strategic play are identical. We also show that the competitive ratio of the greedy strategy is at most 4 (Theorem 5.1).

For the case of empty counters, when players choose undominated strategies, the competitive ratio is unbounded (Theorem 5.2). The greedy strategy isn't well-defined in the setting with empty counters, though players playing greedily with respect to the initial values of resources has similarly poor performance.

We then consider (α, β, γ) -accurate counters. Under greedy strategies, for any non-increasing resource valuation function v_r , the competitive ratio is $O(\alpha^2 \beta)$ (Theorem 5.5) with probability at least $1 - \gamma$. Under undominated strategies, the competitive ratio is bounded only when the function v_r decreases slowly (Theorem 5.8). For arbitrary non-increasing functions, the competitive ratio of undominated strategies is unbounded (Theorem 5.7).

5.1. Perfect counters and empty counters

Before delving into our main results, we point out that, with perfect counters, greedy is the only undominated strategy, and the competitive ratio of greedy is a constant. We state this result formally, and defer its proof to Appendix A.

THEOREM 5.1. *With perfect counters, in the discrete, future-independent setting, greedy behavior is dominant-strategy and all other behavior is strictly dominated. Moreover, $CR_{\text{GREEDY}}(\mathcal{M}_{\text{Full}}, g) = 2$ for each g where each player selects exactly one resource, and $CR_{\text{GREEDY}}(\mathcal{M}_{\text{Full}}, g) = 4$ for any sequential resource-sharing game g .*

Recall, from our example in the introduction, that both greedy and undominated strategies can perform poorly with respect to empty counters. We defer the proof of the following results to Appendix A. Recall that \mathcal{M}_\emptyset refers to the empty counter mechanism.

THEOREM 5.2. *There exist games g such that $CR_{\text{US}}(\mathcal{M}_\emptyset, g)$ is unbounded.*

We show, in the case of private counters, that a restricted class of valuation curves allow for a competitive ratio which is bounded and independent of n for undominated strategies. The next result shows that, for empty counters, even this restriction is not enough to get a competitive ratio which is independent of the number of players n .

THEOREM 5.3. *There exists g such that $CR_{\text{US}}(\mathcal{M}_\emptyset, g) \geq \frac{n^2}{\log(n)}$, when $v_r(t) = \frac{v_r^{\text{init}}}{t}$.*

5.2. Privacy-preserving information: the Greedy Strategy

In this and the following section, we explore the generalizations of Theorem 5.1, to the setting where the counters are not perfect but instead are privacy-preserving. We note that, in the case of perfect counters, the set of undominated strategies contains only the greedy strategy. On the other hand, with incomplete information, both solution concepts are worthy of study. This section extends the result of Section 5.1, showing that approximate counts of the number of people having chosen each resource are sufficient for greedy behavior to approximate OPT . In contrast, in Section 5.3, we show that for arbitrary value curves, undominated strategies perform poorly; for “well-behaved” value curves (such as $v_r(k) = v_r^{\text{init}}/k$), we show undominated strategies will also perform well.

Let a counter be called an *underestimator* if the value of the counter is always (weakly) smaller than the true value it is counting. We mention the following, which allows us to convert arbitrary counters to underestimators. Let the *perceived value* of a player for an action (w.r.t counters) be the value they would get for that action if the counters were perfect.

OBSERVATION 5.4. *An (α, β) -counter vector can be converted to an $(\alpha^2, 2\beta)$ -counter vector which is an underestimator.*

THEOREM 5.5. *Suppose that \mathcal{M} is a (α, β, γ) -counter vector, and that \mathcal{M} is an underestimator. Then, for any discrete, future-independent resource-sharing game g , with probability $1 - \gamma$, $CR_{\text{GREEDY}}(\mathcal{M}, g) = O(\alpha\beta)$.*

LEMMA 5.6. *Suppose players choose greedily according to a (α, β) -underestimator. Then, in sum, their actual value is at least a $\frac{1}{2\alpha\beta}$ -fraction of their perceived value.*

By Theorem 5.1 and Lemma 5.6, a factor dependent only on α, β is lost from OPT when players select their resources greedily. It remains to show that, when players are choosing greedily according to approximate counts, their actual utility is well-approximated by their perceived utility. The proof of the main theorem for this section is below.

PROOF THEOREM 5.5. Assume none of the counters fail which happens with probability at least $1 - \gamma$. Then, let PSW denote the perceived social welfare of a particular

action set, and $OPT_{Counters}$ be the optimal allocation if the displayed value of the counters was correct at each timestep. Then, Lemma 5.6 states that

$$SW(Greedy_{Counters}) \geq \frac{PSW(Greedy_{Counters})}{2\alpha\beta} \quad (1)$$

Thus, we have

$$SW(Greedy_{Counters}) \geq \frac{PSW(Greedy_{Counters})}{2\alpha\beta} \geq \frac{PSW(OPT_{Counters})}{8\alpha\beta} \geq \frac{SW(OPT_{Real})}{8\alpha\beta}$$

where the first inequality comes from 1, the second comes from Theorem 5.1, and the final inequality comes from the fact that the social welfare of OPT on the perceived resource values is at least as high as the social welfare of OPT on the real values, since the counters always under-count. We get this ratio with probability at least $1 - \gamma$, implying the desired bound. \square

PROOF LEMMA 5.6. Suppose k players chose a given resource r . Without loss of generality, let us assume that those k players chose r one after another, for ease of notation: that is, that we are looking at the first k players' counter values. We wish to bound the ratio

$$\frac{\sum_{i=1}^k v_r(s_{i,r})}{\sum_{c=1}^k v_r(c)}.$$

We start by “grouping” the counter values: it cannot take on values that are small for more than a certain number of steps. In particular, if $x_{i,r} > T\alpha\beta$, for some $T \in \mathbb{N}$,

$$s_{i,r} \geq \frac{1}{\alpha}x_{i,r} - \beta \geq \frac{T\alpha\beta}{\alpha} - \beta = (T-1)\beta$$

Now, we bound the ratio from above using this fact.

$$\frac{\sum_{i=1}^k v_r(s_{i,r})}{\sum_{c=1}^k v_r(c)} \leq \frac{2\alpha\beta \sum_{T=1}^{\lceil \frac{k}{\alpha\beta} \rceil} v_r((T-1)\beta)}{\sum_{c=1}^k v_r(c)} \leq \frac{2\alpha\beta \sum_{T=1}^{\lceil \frac{k}{\alpha\beta} \rceil} v_r((T-1)\beta)}{\sum_{T=1}^{\lceil \frac{k}{\alpha\beta} \rceil} v_r((T-1)\beta)} \leq 2\alpha\beta$$

where the first inequality came from the fact that the value curves are non-increasing and the lower bound on the counter values from above, and the second inequality from the fact that all terms are non-negative. \square

5.3. Privacy-preserving counters and Undominated strategies

We begin with an illustration of how undominated strategies can perform poorly for arbitrary value curves, as motivation for the restricted class of value curves we consider in Theorem 5.8. Suppose one resource r has value curve $v_r(1) = H$ and $v_r(i) = 0$ for all $i > 1$, and one resource r' such that $v_{r'}(0) = \epsilon$, and $v_{r'}(i) = 0$ for all $i > 0$. Suppose the second player has both r, r' available. Then it will be undominated for her to choose r' so long as the counters have any possibility of showing her $s_{1,r} < x_{1,r}$. If then the first player has only resource r' available, the total welfare will be just ϵ rather than $H + \epsilon$; this ratio is unbounded.

In the case of greedy players, we were able to avoid the problem of players undervaluing resources rather easily, by forcing the counters to only underestimate $x_{i,r}$. This won't work for undominated strategies: players might assume the counts are shaded downward. While we can bias the counters to only weakly overestimate $x_{i,r}$, if however, the counters include in their possible range the actual value and some value larger than that, any overestimation of $x_{1,r}$ will still imply that r' is undominated for

the second player. More generally, no (ϵ, δ) -privacy-preserving announcement mechanism will be able to differentiate between the case where player 1 chose resource r and where he chose r' with probability much more than $(\epsilon + \delta)$.

THEOREM 5.7. *For an (ϵ, δ) -differentially private announcement mechanism \mathcal{M} , there exist games g for which $CR_{US}(g, \mathcal{M})$ is unbounded.*

Given the above example, we cannot hope to have a theorem as general as Theorem 5.5 when analyzing undominated strategies with privacy-preserving counters. Instead, we show that, for a class of “well-behaved” value curves, we can bound the competitive ratio of undominated strategies (Theorem 5.8).

Again, along the lines of the greedy case, we show that any player who chooses any undominated resource r' over resource r gets a reasonable fraction of the utility she would get from choosing r . Then, by the analysis of greedy players, we have an analogous argument implying the bound of Theorem 5.8.

THEOREM 5.8. *If each value curve v_r has the property that $\psi(\alpha, \beta)v_r(x) \geq v_r(\max\{0, \frac{x}{\alpha^2} - \frac{2\beta}{\alpha}\})$ and also $v_r(\alpha^2x + 2\alpha\beta) \geq \phi(\alpha, \beta)v_r(x)$, then an action profile a of undominated strategies according to (α, β) -counter vector gets an $\left(\frac{1}{\psi(\alpha, \beta)\phi(\alpha, \beta)}\right)$ -fraction of the $SW(\text{GREEDY})$. Thus, $CR_{US}(g) \leq 4\psi(\alpha, \beta)\phi(\alpha, \beta)$.*

In particular, Theorem 5.8 shows that, for games where $v_r(i) = \frac{v_r}{g_r(x_{i,r})}$, where g_r is a polynomial, the competitive ratio of undominated strategies degrades gracefully as a function of the maximum degree of those polynomials. A simple calculation implies the following.

COROLLARY 5.9. *Suppose for a resource-sharing game g , each resource r has a value curve of the form $v_r(x) = \frac{v_r^{init}}{g_r(x)}$, where g_r is a monotonically increasing degree- d polynomial and v_r^{init} is some constant. Then, $CR_{US}(g, \mathcal{M}) \leq O(2\alpha^3\beta)^d$ with \mathcal{M} providing (α, β) -counters.*

We relegate the proof of Theorem 5.8 Appendix B for lack of space.

6. FUTURE-DEPENDENT

The second model of utility we consider is where each player who selects a given resource incurs the same benefit (or cost) from that resource, which is a function of the total weight placed on that resource. That is, player i 's utility for action a_i will be a function of $(\sum_{j=1}^n a_{j,1}, \dots, \sum_{j=1}^n a_{j,m})$ (the total weight placed by all players on resources). We call this setting the *future-dependent* setting. In the future-dependent setting,

$$u_i(a_1, \dots, a_n) = \sum_{r=1}^m a_i^r v_r'(x_r),$$

where we recall that $x_r = \sum_{i'=1}^n a_{i',r}$ is the total utilization of resource r by all players. In this general setting, for many of our results, we will be interested in value curves that do not decrease too quickly. A curve v_r is (w, l) -shallow if

$$\max_{x \leq l} \frac{\int_0^x v_r'(t) dt}{xv_r'(x)} \leq w$$

which says that the integral of v_r' from 0 to x (what players would see as their payoff from resource r as each made a decision) is not too much larger than the the actual payoff all players get from the resource being utilized with x weight.

In Section C, we show that this restriction on the *rate of decay* of the value curves is necessary to say anything nontrivial about the performance of the greedy strategy. We show that, even with perfect counters, greedy's performance is $\Theta(w)$, where all of the curves are guaranteed to be (w, n) -shallow. Due to space constraints, we focus on the special case of market sharing in the section below.

6.1. Market sharing

Market sharing is the special case of $v_r(x_r) = c/x_r$ for all $x_r \geq 1$. Goemans et al. [2004] showed that for market-sharing games, the competitive ratio of α -approximate greedy play is at most $O(\alpha \log(n))$. Using analysis similar to theirs, we have the following result.

COROLLARY 6.1. *With (α, β, γ) -counter vector and greedy play, with probability at least $1 - \gamma$, the welfare achieved is at least $(OPT - 2\beta\alpha n)/O((1 + \alpha^2) \log(n))$.*

We now focus our attention on a game play based on undominated strategy.

THEOREM 6.2. *With perfect counters and undominated strategic play, there are games for which the welfare achieved is at most $OPT/(n \log(n))$.*

PROOF. Here is an example with n players. Consider the case where for every $i \geq 1$, player i is interested in resource 0 and resource i . For every $i \geq 1$, the total value of resource i is $(n - i + 1)(1 - \epsilon)/i$ (for some small $\epsilon > 0$). The value of resource 0 is 1.

We claim that there is undominated strategy game play where every player chooses resource 0 giving a social welfare of 1, whereas the optimal welfare is achieved by assigning player i resource i giving a total welfare of $n(\log(n) - 1)(1 - \epsilon)$.

Here is such an undominated strategy profile: for each i , player i believes that every player after her is only interested in resource i . With this belief, it is easy to see that choosing resource 0 is an undominated strategy for every player. \square

7. PRIVATE COUNTER VECTORS WITH LOWER ERRORS FOR SMALL VALUES

In this section, we describe a counter for the model of differential privacy under continual observation that has improved guarantees when the value of the counter is small.

Recall the basic counter problem: given a stream $\vec{a} = (a_1, a_2, \dots, a_n)$ of numbers $a_i \in [0, 1]$, we wish to release at every time step t the partial sum $x_t = \sum_{i=1}^t a_i$.

We require a generalization, where one maintains a vector of m counters. Each player's update contribution is now a vector $a_i \in [0, 1]^m$, with the constraint that $\|a_i\|_1 \leq 1$. That is, a player can add non-negative values to all counters, but the total value of her updates is at most 1. The partial sums x_t then lie in $(\mathbb{R}^+)^m$ (with ℓ_1 norm bounded by t).

Given an algorithm \mathcal{A} , we define the output stream $(s_1, s_2, \dots, s_n) = \mathcal{A}(\vec{a})$ where $s_i = \mathcal{A}(t, a_1, \dots, a_{i-1})$. The original works on differentially private counters [Dwork et al. 2010; Chan et al. 2011] concentrated on minimizing the additive error of the estimated sums, that is, they sought to minimize $\|x_t - s_t\|_\infty$. Both papers gave a binary tree-based mechanism, which we dub "TreeSum", with additive error approximately $(\log^2 n)/\epsilon$. Some of our algorithms use TreeSum, and others use a new mechanism (FTSum, described below) which gets a better additive error guarantee at the price of introducing a small multiplicative error. We capture a mixed approximation guarantee as follows:

Definition 7.1. The algorithm \mathcal{A} provides an (α, β, γ) -approximation to partial sums if for every (adaptively defined) sequence $\vec{a} \in ([0, 1]^m)^n$, with probability at least $1 - \gamma$ over the coins of \mathcal{A} , for all times $i \in [n]$ and counters $r \in [m]$, the reported value $x_{t,r}$

satisfies:

$$\frac{1}{\alpha} \cdot x_{i,r} - \beta \leq s_{i,r} \leq \alpha \cdot x_{i,r} + \beta.$$

Proofs of all the results in this section can be found in Appendix E.

LEMMA 7.2. *For every $m \in \mathbb{N}$ and $\gamma \in (0, 1)$: Running m independent copies of TreeSum [Dwork et al. 2010; Chan et al. 2011] is $(\epsilon, 0)$ -differentially private and provides an $(1, C_{tree} \cdot \frac{(\log n)(\log(nm/\gamma))}{\epsilon}, \gamma)$ -approximation to partial vector sums, where $C_{tree} > 0$ is an absolute constant.*

Even for $m = 1, \alpha = 1$, this bound is slightly tighter than those in Chan et al. [2011] and Dwork et al. [2010]; however, it follows directly from the tail bound in Chan et al. [2011].

Our new algorithm, FTSum (for Flag/Tree Sum), is described in Algorithm 1. For small m ($m = o(\log(n))$), it provides lower additive error at the expense of introducing an arbitrarily small constant multiplicative error.

LEMMA 7.3. *For every $m \in \mathbb{N}$, $\alpha > 1$ and $\gamma \in (0, 1)$, FTSum (Algorithm 1) is $(\epsilon, 0)$ -differentially private and $(\alpha, \tilde{O}_\alpha(\frac{m \log(n/\gamma)}{\epsilon}), \gamma)$ -approximates partial sums (where $\tilde{O}_\alpha(\cdot)$ hides polylogarithmic factors in its argument, and treats α as constant).*

FTSum proceeds in two phases. In the first phase, it increments the reported output value only when the underlying counter value has increased significantly. Specifically, the mechanism outputs a public signal, which we will call a “flag”, roughly when the true counter achieves the values $\log n$, $\alpha \log n$, $\alpha^2 \log n$ and so on, where α is the desired *multiplicative* approximation. The reported estimate is updated each time a flag is raised (it starts at 0, and then increases to $\log n$, $\alpha \log n$, etc). The privacy analysis for this phase is based on the “sparse vector” technique of Hardt and Rothblum [2010], which shows that the cost to privacy is proportional to the number of times a flag is raised (but not the number of time steps between flags).

When the value of the counter becomes large (about $\frac{\alpha \log^2 n}{(\alpha-1)\epsilon}$), the algorithm switches to the second phase and simply uses the TreeSum protocol, whose additive error (about $\frac{\log^2 n}{\epsilon}$) is low enough to provide an α multiplicative guarantee (without need for the extra space given by the additive approximation).

If the mechanism were to raise a flag *exactly* when the true counter achieved the values $\log n$, $\alpha \log n$, $\alpha^2 \log n$, etc, then the mechanism would provide a $(\alpha, \log n, 0)$ approximation during the first phase, and a $(\alpha, 0, 0)$ approximation thereafter. The rigorous analysis is more complicated, since flags are raised only near those thresholds.

PROPOSITION 7.4. *If \mathcal{A} is (ϵ, δ) -private and (α, β, γ) -accurate, then one can modify \mathcal{A} to obtain an algorithm \mathcal{A}' with the same efficiency that is $(\epsilon, \delta + \gamma)$ -private and $(\alpha, \beta, 0)$ -accurate.*

COROLLARY 7.5. *Algorithm 1 is an (ϵ, δ) -differentially private vector counter algorithm providing a*

- (1) $(1, O(\frac{(\log n)(\log(nm/\delta))}{\epsilon}), 0)$ -approximation (using modified TreeSum); or
- (2) $(\alpha, \tilde{O}_\alpha(\frac{m \log n \log \log(1/\delta)}{\epsilon}), 0)$ -approximation for any constant $\alpha > 1$ (using FTSum).

8. OTHER GAMES

In this section, we study a number of games which Leme et al. [2012] showed to have a large improvement between their Price of Anarchy and their *sequential* Price of Anarchy. We pose the question: with privacy-preserving information handed out to players,

Algorithm 1: FTSum — A Private Counter with Low Multiplicative Error**Input:** Stream $\vec{a} = (a_1, \dots, a_n) \in ([0, 1]^m)^n$, parameters $m, n \in \mathbb{N}$, $\alpha > 1$ and $\gamma > 0$ **Output:** Noisy partial sums $s_1, \dots, s_n \in \mathbb{R}^m$ $k \leftarrow \lceil \log_\alpha \left(\frac{\alpha}{\alpha-1} \cdot C_{tree} \cdot \frac{\log(nm/\gamma)}{\epsilon} \right) \rceil$;/* C_{tree} is the constant from Lemma 7.2 */ $\epsilon' \leftarrow \frac{\epsilon}{2m(k+1)}$;**for** $r = 1$ **to** m **do** $\text{flag}_r \leftarrow 0$; $x_{0,r} \leftarrow 0$; $\tau_r \leftarrow (\log n) + \text{Lap}(2/\epsilon')$;**for** $i = 1$ **to** n **do** **for** $r = 1$ **to** m **do** **if** $\text{flag}_r \leq k$ **then** (First phase still in progress for counter r) $x_{i,r} \leftarrow x_{i-1,r} + a_{i,r}$; $\tilde{x}_{i,r} \leftarrow x_{i,r} + \text{Lap}(\frac{2}{\epsilon'})$; **if** $\tilde{x}_{i,r} > \tau_r$ **then** (Raise a new flag for counter r) $\text{flag}_r \leftarrow \text{flag}_r + 1$; $\tau_r \leftarrow (\log n) \cdot \alpha^{\text{flag}_r} + \text{Lap}(2/\epsilon')$; **Release** $s_{i,r} = (\log n) \cdot \alpha^{\text{flag}_r - 1}$; **else** (Second phase has been reached for counter r) **Release** $s_{i,r} = r$ -th counter output from $\text{TreeSum}(\vec{a}, \epsilon/2)$;

what loss is incurred in comparison to providing exact information? In addition to introducing privacy constraints, we should note here that while in Leme et al. [2012], each player playing the sequential game knows the *type of every other player*, in our setting, we only provide information about *actions taken by previous players*. For most results in this section, we relegate the proof to Appendix F.

8.1. Unrelated Machine scheduling games

An instance of the unrelated machine scheduling game consists of n players who must schedule their respective jobs on one of the m machines; the cost to the player is the final load on the machine on which she scheduled her job. The size of player k 's job on machine q is t_{kq} . The objective of the mechanism is to minimize the makespan. We consider the dynamics of this game when it is played sequentially. Leme et al. [2012] prove that the sequential price of anarchy is $O(m2^n)$.

In our setting, each player is shown a *load profile* when it is her turn to play. The *load profile* L denotes the *displayed* vector of loads on the various machines. In the *perfect counter* setting, the displayed load equals the exact load on each machine. We now show that undominated strategies, with perfect counters, perform unboundedly poorly with respect to OPT.

LEMMA 8.1. *If \mathcal{M} is a perfect counter vector, $CR_{US}(\mathcal{M}, g)$ is unbounded for some instances g of unrelated machine scheduling.*

PROOF. Consider the case with two players (p_1 and p_2) and two machines (m_1 and m_2). p_1 arrives before p_2 . Player p_1 has a cost of 0 on m_1 and 1 on m_2 . It is an undominated strategy for player 1 to choose m_2 since if player p_2 has a cost of 2 on m_1 and 3 on m_2 , p_2 chooses m_1 and so p_1 is better off scheduling her job on m_2 .

However, if p_1 chooses m_2 (an undominated strategy) m_2 , and player p_2 has cost 1 on m_1 and 0 on m_2 , the optimal makespan is 0; the achieved makespan is at least 1. \square

In light of this result, we restrict our attention to greedy strategies for machine scheduling, and show that the competitive ratio of the greedy strategy with privacy-preserving counters is bounded. Below, we denote by t_k^* the minimum cost of job k among all the machines. The following result follows from the analysis of the greedy algorithm as presented in Aspnes et al. [1997].

THEOREM 8.2. *[Aspnes et al. 1997] With perfect counters and players playing greedy strategies, the makespan is at most $\sum_{i=1}^n t_i^*$, and since $OPT \geq \sum_{i=1}^n t_i^*/m$, the competitive ratio is at most m .*

Theorem 8.3 shows that such a bound extends to the setting where players have only approximate information about the state, showing that privacy-preserving information is enough to attain nontrivial coordination with greedy players.

THEOREM 8.3. *Using (α, β, γ) -counter vector, and players playing greedy strategies, with probability $1 - \gamma$, the makespan is at most $\alpha^{2n+1}m \cdot OPT + \beta(\alpha^{2n+1}(2n+1) + 1)$.*

8.2. Cut games

A cut game is defined by a graph, where every player is a node of the graph. Each of the n players chooses one of the two colors, ‘red’ or ‘blue’, and the utility to a player is the number of her neighbors who *do not have* the same color as hers.

In sequential play, when a player has her turn to play, she is shown counts of the number of her neighbors who are colored ‘red’ and who are colored ‘blue’. We assume each player knows the total number of her neighbors in the graph exactly. With greedy strategies, each player chooses the color with fewer nodes when it is her turn to play. As was the case for machine scheduling, undominated strategies for cut games perform much worse than OPT , even with perfect counters.

LEMMA 8.4. *With perfect counters and undominated strategies, the competitive ratio against the optimal social welfare is at least n .*

Given the previous result, we focus our attention on greedy strategies. With greedy strategies and perfect counters, the competitive ratio is constant, shown by Leme et al. [2012]. We show that, with privacy-preserving counters, it is possible to compare the social welfare of greedy to that of OPT .

THEOREM 8.5. *[Leme et al. 2012] With perfect counters and greedy strategies, the competitive ratio against the optimal social welfare is at most 2.*

Now, we compare the performance of greedy w.r.t. to approximate counters to OPT .

THEOREM 8.6. *With (α, β, γ) -counter vector and greedy strategies, with probability at least $1 - \gamma$, the social welfare is at least $\frac{OPT}{2\alpha^2} - \frac{2\beta}{\alpha}n$.*

8.3. Cost sharing games

A cost sharing game is defined as follows. n players each have to choose one of the m sets. There is an underlying bipartite graph between the players and the sets, and a player can choose only one among those sets that she is adjacent to (i.e., she shares an edge with). Moreover, every set i has a cost c_i and the cost to a player is the cost of the set she chooses divided by the number of players who chose that set i.e., each of the players who choose a particular set share its cost equally. Each player would like to minimize her cost; the social welfare is the sum of costs of the players, which is equal to the sum of the costs of the sets chosen by various players.

Leme et al. [2012] prove that the sequential price of anarchy is $O(\log(n))$. Our work uses counters to publicly display an estimate of the number of the players who have

selected that set so far. With perfect counters, this estimate is always exact. Unfortunately, greedy strategies can perform poorly in this setting, even with exact counters.

LEMMA 8.7. *With perfect counters and greedy strategies, the competitive ratio is n .*

PROOF. We first show that the competitive ratio is at most n . Let \hat{s}_i be the set that i should choose in the optimal allocation, and let s_i she chose. Also, let $l(s_i)$ be the number of player who chose set s_i . Greedy strategy dictates that it must be the case that $c_{s_i}/l(s_i) \leq c_{\hat{s}_i}$. Summing over all players i , we have the total cost of the allocation produced by the mechanism is $\sum_{i=1}^n c_{s_i}/l(s_i) \leq \sum_{i=1}^n c_{\hat{s}_i}$, and this is equal to $\sum_{j \in J} q(j)c_j$, where J is the collection of sets picked in the optimal allocation and q_j is the number of players allocated to set j . Since the optimal cost is $\sum_{j \in J} c_j$ and $q_j \leq n$, we have the competitive ratio is at most n .

We now show that the competitive ratio is at least n . Consider the case where there is a public set s that is adjacent to all the players and has cost $1 + \epsilon$ (for any small $\epsilon > 0$). In addition, there are n private sets s_1, \dots, s_n with set s_i having cost 1 and adjacent only to player i . In the sequential game play, with greedy strategies and perfect counters (indicating the number of players who have chosen a particular set so far in the game), each player will choose her private set since that will have cost 1 as opposed to $1 + \epsilon$ for the public set. This gives a total cost of n . The optimal solution is to pick the public set with a total cost of $1 + \epsilon$. \square

In light of Lemma 8.7, the greedy strategy with respect to approximate counters should not perform well with respect to *OPT*. However, we do show that there are instances in which greedy with respect to these approximate counters can be better than greedy with respect to perfect counters. The example we use is the same as in Lemma 8.7, and is also to the example showing the price of anarchy for cost-sharing is $\Omega(n)$. Proposition 8.8 and the exponential improvement of the sequential price of anarchy over the simultaneous price of anarchy [Leme et al. 2012] suggest the instability of this equilibrium.

PROPOSITION 8.8. *In certain instances of cost sharing with greedy strategies, the competitive ratio using privacy-preserving counters is better than using perfect counters.*

PROOF. Consider the same instance as in Lemma 8.7. There is a public set that is adjacent to all the players and has cost $1 + \epsilon$. In addition, there is a private set for each player that is adjacent to only that player. Each private set has cost 1. The number of players is n and the number of sets is $m = n + 1$.

Consider the following construction of the counter vector (here $p = 1$, $q = O(\log(n) \log(n^2 m)/\epsilon)$, $r = 1/n$ and $c = 8(p^2 + 2pq)$). For the initial sequence of c players, for each player $i \in [c]$, for each counter, a uniformly randomly chosen number in the range $[0, c]$ (drawn independently for each counter) is displayed. Starting with the $(c + 1)$ st player, each counter in the counter vector displays the value according to (p, q, r) -Tree-sum based construction (Lemma 7.2). It is easy to verify that the construction gives a (α, β, γ) counter vector for $\alpha = p$, $\beta = c$ and $\gamma = r$.

Let P be the counter that corresponds to the public set, and S_i be the counter for the i^{th} private set in the counter vector. Initially, the true value of all the counters is 0. For the initial set of c players, for each $i \in [c]$, the probability that the displayed value of P is greater than that of S_i is $1/2$ (since for each player $i \in [c]$, on each counter, a uniformly random number drawn independently from the range $[0, c]$ is displayed).

Hence, in the first c players, the expected number of players for who the displayed value of P is greater than the corresponding S_i is $c/2$, and under greedy strategy, all these players will choose the public set. Hence the expected true count of the P at the

end of the prefix of c players is $c/2$. Using a Chernoff bound, the probability the true count of P after the first c players is smaller than $c/4$ is at most $e^{(-c/16)}$.

After the initial sequence of c players, the counter values are displayed according to the (p, q, r) -Tree based construction. By the error guarantees, it follows that if the true count for the public set is at least $p^2 + 2pq$ at the end of the initial c -length sequence, then for the rest of the players, with probability $(1 - r)$, the displayed value of P is *always strictly greater* than the displayed value of every S_i (whose true value is at most 1 and so the displayed value is at most $p + q$). Since $c/4 = 2(p^2 + 2pq)$, we can infer that with probability at least $(1 - r - e^{(-c/16)})$, all players after the initial sequence of length c will choose the public set giving the total cost of at most $1 + \epsilon + c$. In contrast, with perfect counters, the total cost is always n (Lemma 8.7). \square

9. DISCUSSION AND OPEN PROBLEMS

In this work, we considered how public dissemination of information in sequential games can guarantee a good social welfare while maintaining differential privacy of the players' strategies. We considered two 'extreme' cases – the greedy strategy and the class of all undominated strategies. While analyzing the class of undominated strategies gives guarantees that are robust, in many games that we considered, the competitive ratios were significantly worse than greedy strategies, and in some cases they were unbounded. It is an interesting direction for future research to consider classes of strategies that more restricted than undominated strategies yet are general enough to be relevant for games where players play with imperfect information.

As mentioned in the introduction, we note here that, while players are making choices subject to approximate information, our results are not a direct extension of the line of thought that approximate information implies approximate optimization. In particular, for greedy strategies, while there may be a bound on the error of the counters, but that *does not imply* that, for arbitrary value curves, playing greedily according to the counters will be *approximately optimal for each individual*. In particular, consider one resource r with value H for the first 10 investors, and value 0 for the remaining investors, and a second resource r' with value $H/2$ for all investors. With (α, β, γ) , as many as β players might have unbounded ratio between their value for r as r' , but will pick r over r' . The analysis of greedy shows, despite this anomaly, the total social welfare is still well-approximated by this behavior.

All of our results relied on using differentially private counters for disseminating information. For the differentially-private counter, a main open question is "what is the optimal trade-off between additive and multiplicative guarantees?". Furthermore, as part of future research, one can consider other privacy techniques for announcing information that can prove useful in helping players achieve a good social welfare. And more generally, we want to understand what features of games lend themselves to be amenable to public dissemination of information that helps achieve good welfare and simultaneously preserves privacy of the players' strategies.

REFERENCES

- ASPINES, J., AZAR, Y., FIAT, A., PLOTKIN, S., AND WAARTS, O. 1997. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *J. ACM* 44, 3, 486–504.
- AWERBUCH, B., AZAR, Y., GROVE, E. F., KAO, M.-Y., KRISHNAN, P., AND VITTER, J. S. 1995. Load balancing in the l_p norm. In *FOCS '95*. IEEE, 383–391.
- BALCAN, M.-F., BLUM, A., AND MANSOUR, Y. 2009. The price of uncertainty. *EC '09*. 285–294.
- CHAN, T.-H. H., SHI, E., AND SONG, D. 2011. Private and continual release of statistics. *ACM Trans. Inf. Syst. Secur.* 14, 3, 26.
- CHEN, Y., CHONG, S., KASH, I. A., MORAN, T., AND VADHAN, S. 2013. Truthful mechanisms for agents that value privacy. *EC '13*. 215–232.

- CHRISTODOULOU, G. AND KOUTSOUPIS, E. 2005a. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Algorithms-ESA 2005*. Springer, 59–70.
- CHRISTODOULOU, G. AND KOUTSOUPIS, E. 2005b. The price of anarchy of finite congestion games. In *STOC '05*. ACM, 67–73.
- DWORK, C., MCSHERRY, F., NISSIM, K., AND SMITH, A. 2006. Calibrating noise to sensitivity in private data analysis. *TCC'06*. Springer-Verlag, Berlin, Heidelberg, 265–284.
- DWORK, C., NAOR, M., PITASSI, T., AND ROTHBLUM, G. N. 2010. Differential privacy under continual observation. *STOC '10*. ACM, 715–724.
- FINANCIAL CRISIS INQUIRY COMMISSION. 2011. *The Financial Crisis Inquiry Report: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the United States*. U.S. Government Printing Office.
- FLEISCHER, L. K. AND LYU, Y.-H. 2012. Approximately optimal auctions for selling privacy when costs are correlated with data. *EC '12*. ACM, 568–585.
- FLOOD, M. D., KATZ, J., ONG, S. J., AND SMITH, A. 2013. Cryptography and the economics of supervisory information: Balancing transparency and confidentiality. Working Paper #11, Office of Financial Research, US Department of Treasury. August.
- GHOSH, A. AND ROTH, A. 2011. Selling privacy at auction. *EC '11*. ACM, 199–208.
- GOEMANS, M., LI, L. E., MIRROKNI, V. S., AND THOTTAN, M. 2004. Market sharing games applied to content distribution in ad-hoc networks. *MobiHoc '04*. ACM, 55–66.
- HARDT, M. AND ROTHBLUM, G. N. 2010. A multiplicative weights mechanism for privacy-preserving data analysis. In *FOCS '10*.
- HSU, J., HUANG, Z., ROTH, A., ROUGHGARDEN, T., AND WU, Z. S. 2013. Private matchings and allocations. *CoRR abs/1311.2828*.
- HUANG, Z. AND KANNAN, S. 2012. The exponential mechanism for social welfare: Private, truthful, and nearly optimal. *FOCS '12*. 140–149.
- KARP, R. M., VAZIRANI, U. V., AND VAZIRANI, V. V. 1990. An optimal algorithm for on-line bipartite matching. In *STOC '90*. 352–358.
- KEARNS, M., PAI, M. M., ROTH, A., AND ULLMAN, J. 2012. Mechanism design in large games: Incentives and privacy. *CoRR abs/1207.4084*.
- LEME, R. P., SYRGKANIS, V., AND TARDOS, E. 2012. The curse of simultaneity. *ITCS '12*. 60–67.
- LIBMAN, L. AND ORDA, A. 1997. Atomic resource sharing in noncooperative networks. *Telecommunication Systems* 17, 2001.
- LIGETT, K. AND ROTH, A. 2012. Take it or leave it: Running a survey when privacy comes at a cost. In *WINE' 12*. Vol. 7695. 378–391.
- MCSHERRY, F. AND TALWAR, K. 2007. Mechanism design via differential privacy. *FOCS '07*. 94–103.
- NISSIM, K., ORLANDI, C., AND SMORODINSKY, R. 2012a. Privacy-aware mechanism design. *EC '12*. ACM, 774–789.
- NISSIM, K., SMORODINSKY, R., AND TENNENHOLTZ, M. 2012b. Approximately optimal mechanism design via differential privacy. *ITCS '12*. ACM, 203–213.
- OET, M. V., BIANCO, T., GRAMLICH, D., AND ONG, S. J. 2011. Safe: An early warning system for systemic banking risk. Working Paper 11-29, Federal Reserve Bank of Cleveland.
- OET, M. V., BIANCO, T., GRAMLICH, D., AND ONG, S. J. 2012. Financial stress index: A lens for supervising the financial system. Working Paper 12-37, Federal Reserve Bank of Cleveland.
- PAI, M. M. AND ROTH, A. 2013. Privacy and mechanism design. *CoRR abs/1306.2083*.
- ROGERS, R. AND ROTH, A. 2013. Asymptotically truthful equilibrium selection in large congestion games. *CoRR abs/1311.2625*.
- ROTH, A. 2011. Cis 800/002, fall 2011: The algorithmic foundations of data privacy.
- ROTH, A. AND SCHOENEBECK, G. 2012. Conducting truthful surveys, cheaply. *EC '12*. ACM, 826–843.
- ROUGHGARDEN, T. 2003. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences* 67, 2, 341–364.
- ROUGHGARDEN, T. AND TARDOS, É. 2004. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior* 47, 2, 389–403.
- SURI, S., TÓTH, C. D., AND ZHOU, Y. 2007. Selfish load balancing and atomic congestion games. *Algorithmica* 47, 1, 79–96.
- XIAO, D. 2013. Is privacy compatible with truthfulness? *ITCS '13*. ACM, 67–86.

Online Appendix to: Privacy-Preserving Public Information for Sequential Games

AVRIM BLUM, Carnegie Mellon University
JAMIE MORGENSTERN, Carnegie Mellon University
ANKIT SHARMA, Carnegie Mellon University
ADAM SMITH, Penn State University

A. FUTURE INDEPENDENT: DISCRETE VERSION

THEOREM A.1. *With perfect counters, in the discrete, future-independent setting, greedy behavior is dominant-strategy and all other behavior is strictly dominated. Moreover, $CR_{\text{GREEDY}}(\mathcal{M}_{\text{Full}}, g) = 2$ for each g where each player selects exactly one resource, and $CR_{\text{GREEDY}}(\mathcal{M}_{\text{Full}}, g) = 4$ for any sequential resource-sharing game g .*

The proof of this Theorem follows from the connection between future-independent resource-sharing and online vertex-weighted matching, which we mention below.

OBSERVATION A.2. *In the setting where $\|a_i\|_1 = 1$ for all $a_i \in A_i$, for all i , the full-information, discrete, future-independent setting reduces to online, vertex-weighted bipartite matching. The full-information, discrete, future-independent setting reduces to many-to-one online, vertex-weighted bipartite matching where the vertices arriving online have a set of subsets of allowable matches.*

PROOF. Construct the following bipartite graph $G = (U, V, E)$ as an instance of online vertex-weighted matching from an instance of the future-independent resource sharing game. For each resource r , create n vertices in V , one with weight $v_r(t)$ for each $t \in [n]$. As players arrive online, they will correspond to vertices in $u_i \in U$. For each $a_i \in A_i$ corresponding to a set of resources S , u_i is allowed to take any subset of V with a single copy of each $r \in S$. \square

LEMMA A.3 ([KARP ET AL. 1990]). *The greedy strategy for online, vertex-weighted bipartite matching has a competitive ratio of $\frac{1}{2}$.*

PROOF OF THEOREM 5.1. Consider any instance of $G = (U, V, E)$, a vertex-weighted bipartite graph. Let μ be the optimal many-to-one matching from U to V (where $u \in U$ has potentially multiple neighbors in V). Consider μ' , the greedy many-to-one matching for a particular sequence of arrivals σ .

Consider a particular $u \in U$, and the time it arrives $\sigma(u)$ as μ' progresses. If at least $1/2$ the value of $\mu(u)$ is available at that time, then $w(\mu'(u)) \geq \frac{1}{2}w(\mu(u))$ (since u can be matched to any subset of $\mu(u)$, by the downward closed assumption). If not, then $w(\mu'(\mu(u))) \geq \frac{1}{2}w(\mu(u))$ (at least half the value was taken by others). Thus, we know that, for all u ,

$$w(\mu'(u)) + w(\mu'(\mu(u))) \geq \frac{1}{2}w(\mu(u))$$

summing up over all u , we get

$$\sum_u w(\mu'(u)) + w(\mu'(\mu(u))) = 2w(\mu') \geq \frac{1}{2} \sum_u w(\mu(u)) = \frac{1}{2}w(\mu)$$

Rearranging shows that $w(\mu') \geq \frac{1}{4}w(\mu)$. \square

Note that this proof also applies when, rather than each vertex having a set of downwards-closed subsets he can choose between, each vertex has a set of edges and can choose any subset of those edges, though that's not necessary for our purposes.

A.1. Empty Counters for Resource Sharing: Lower bounds

In this section, we investigate what happens with empty counters. Suppose players have no knowledge of other players' allowable actions, and no information about other players' selected actions. Let \mathcal{M}_\emptyset denote the mechanism that outputs $\bar{0}$ for all inputs.

PROOF OF THEOREM 5.2. Let g be the following game. For each player i , there is a resource r_i such that $v_{r_i}(1) = H$ but $v_{r_i}(> 1) = 0$. Furthermore, let there be some other resource r such that $v_r(1) = 1$. Let A_i contain 2 allowable actions: selecting r_i and selecting r .

OPT in this setting would have each player select r_i , which has $SW(OPT) = nH$. On the other hand, we claim it is undominated for each player to select r instead (call this joint action a). If each player were to have a "twin", then r_i could have already been selected by another player so that i would get more utility from r than r_i . Then, this undominated strategy a has $SW(a) = n$. Thus, we have a game g for which

$$CR_{US}(g) \geq \frac{nH}{n} = H$$

which, as $H \rightarrow \infty$ is unbounded. \square

The negative result above isn't particularly surprising: if there is some coordination to be done, but there is no coordinator and no information about the target, all is lost. On the other hand, our positive result for undominated strategies (Theorem 5.8) in the case of private information relies on a very particular rate of decay of the resources' value. Theorem 5.3 show that, even under this stylized assumption where all resources' values shrink slowly, a total lack of information can lead to very poor behaviour in undominated strategies.

PROOF OF THEOREM 5.3. For each player i , let r_i be a resource where $v_{r_i}(1) = n$ (note that this uniquely determines $v_i(c)$ for all c). Let there be another resource r such that $v_r(1) = 1$. Let each A_i contain all resources. Since $\frac{v_{r_i}(1)}{n} = 1$, it is not dominated for player i to select r . Let a denote the joint strategy where each player selects resource r . Then, $SW(A) = \log(n)$. Since $SW(OPT) = n^2$, $CR_{US} \geq \frac{n^2}{\log(n)}$. \square

B. UNDOMINATED STRATEGIES FOR RESOURCE-SHARING, OMITTED PROOFS

PROOF OF THEOREM 5.8. Consider a player i . We show that any undominated set of resources R' gets a reasonable fraction of the greedy resource set choice R . For a resource r , there is a conceivable range of $x_{i,r}$. Let $\widehat{x_{i,r}}$ denote player i 's reasoned possible value of $x_{i,r}$, consistent with the announcement $s_{i,r}$. Now, consider the current true count $x_{i,r}$. We will directly argue about the possible range of the perceived counts $\widehat{x_{i,r}}$ as a function of the true count. By the bounds on (α, β, γ) -counters, for a given true value x , it must be the case that all announcements $s_{i,r}$ satisfy:

$$\alpha x_{i,r} + \beta \geq s_{i,r} \geq \frac{1}{\alpha} x_{i,r} - \beta$$

Rearranging, we have $s_{i,r} \in [\frac{1}{\alpha} x_{i,r} - \beta, \alpha x_{i,r} + \beta]$. Suppose these bounds are realized; we wish to upper and lower bound $\widehat{x_{i,r}}$ as a function of these announcement values. By the quality of the announcement, we have that $\alpha \widehat{x_{i,r}} + \beta \geq s_{i,r} \geq \frac{1}{\alpha} x_{i,r} - \beta$.

We can similarly upper bound $\widehat{x_{i,r}}$, e.g. $\alpha x_{i,r} + \beta \geq s_{i,r} \geq \frac{1}{\alpha} \widehat{x_{i,r}} - \beta$, which, by the fact that the true count is at least 0, implies $\widehat{x_{i,r}} \in [\max\{0, \frac{x_{i,r}}{\alpha^2} - \frac{2\beta}{\alpha}\}, \alpha^2 x_{i,r} + 2\alpha\beta]$. Now, suppose player i chose the set of resource R' which was undominated, while R would have been the greedy choice. Since the set R' is undominated (by the greedy choice R):

$$\sum_{r \in R'} v_r(\widehat{x_{i,r}}) \geq \sum_{r \in R} v_r(\widehat{x_{i,r}}) \quad (2)$$

We use the lower bound on $\widehat{x_{i,r}}$ to imply

$$\sum_{r \in R'} v_r(\widehat{x_{i,r}}) \leq \sum_{r \in R'} v_r(\max\{0, \frac{x_{i,r}}{\alpha^2} - \frac{2\beta}{\alpha}\}) \leq \sum_{r \in R'} \psi(\alpha, \beta) v_r(x_{i,r}) \quad (3)$$

where the first inequality came from the lower bound on the counter, and the fact that the valuations are decreasing, and the second from the assumption about v_r on x and its lower bound.

Similarly, we know for each r that

$$\sum_{r \in R} v_r(\widehat{x_{i,r}}) \geq \sum_{r \in R} v_r(\alpha^2 x_{i,r} + 2\alpha\beta) \geq \sum_{r \in R} \frac{v_r(x_{i,r})}{\phi(\alpha, \beta)} \quad (4)$$

Combining (4,3,2)

$$\psi(\alpha, \beta) \phi(\alpha, \beta) \sum_{r \in R'} v_r(x_{i,r}) \geq \sum_{r \in R} v_r(\widehat{x_{i,r}})$$

we have the desired ratio. \square

C. FUTURE-DEPENDENT, GENERAL CASE

THEOREM C.1. *There exist sequential resource-sharing games g , where each resource r 's value curve v_r' is (w, n) – shallow, such that in the full-information, future-dependent setting, $CR_{\text{GREEDY}}(\mathcal{M}_{\text{Full}}, g) \geq 2w$.*

PROOF. Consider two players and two resources r, r' . Let r have a value curve $v_r'(0) = w$, $v_r'(1) = \frac{1}{2}$ and $v_{r'}'(0) = w - \epsilon$. Suppose player one has access to both resources, the other having only resource r as an option. Then, player one will choose r according to greedy, and player two will always select r . The social welfare will be $SW(\text{greedy}) = 1$, whereas OPT is for player 1 to take r' and will have $SW(OPT) = 2w - \epsilon$. As $\epsilon \rightarrow 0$, this ratio approaches $2w$. \square

Thus, as $w \rightarrow \infty$, the competitive ratio of the greedy strategy is unbounded. Fortunately, the competitive ratio cannot be worse than this, for fixed w , as we show in the theorem below.

THEOREM C.2. *Suppose, for a sequential resource-sharing game g , each resource r 's value curve v_r is (w, n) – shallow. Then, in the full-information, future-dependent setting, $CR_{\text{GREEDY}}(\mathcal{M}_{\text{Full}}, g) \leq 4w$.*

PROOF. Let PSW denote the perceived social welfare of a particular action set. Consider a given player who chooses some set of resources according to the greedy strategy, with the impression that she should get value V from her choices. That is, she chose a_i such that

$$a_i = \operatorname{argmax}_{a_i \in A_i} \sum_r \int_{x_{i,r}}^{x_{i,r} + a_{i,r}} v_r'(x) dx$$

Then, we sum up the perceived utility all players have for their actions a_i :

$$\begin{aligned} PSW(\text{GREEDY}) &= \sum_{i \in [n]} \sum_r \int_{x_{i,r}}^{x_{i,r} + a_{i,r}} v_r'(x) dx = \sum_r \int_0^{x_{n,r} + a_{n,r}} v_r'(x) dx \\ &\leq w(x_{n,r} + a_{n,r}) v_r'(x_{n,r} + a_{n,r}) = w SW(\text{GREEDY}) \end{aligned} \quad (5)$$

where the last inequality comes from our assumption about the value curves all being (w, n) – shallow.

We now need to relate this quantity to OPT . Consider the game g' where each player actually received her perceived payoff $\sum_r \int_{x_{i,r}}^{x_{i,r} + a_{i,r}} v_r'(x) dx$. It is the case that $OPT_{g'} \geq OPT_g$. Moreover, players are choosing their strategies greedily according to g' 's utility functions, so

$$SW(\text{GREEDY}, g) \geq \frac{1}{w} PSW(\text{GREEDY}, g) = \frac{1}{w} SW(\text{GREEDY}, g') \geq \frac{1}{4w} OPT_{g'} \geq \frac{1}{4w} OPT_g$$

where the first inequality follows from (5), the second from the fact that $PSW(g, a) = SW(g', a)$ for all a , the third from the fact that GREEDY is 4-competitive with OPT for g' by Lemma 5.1, and the final inequality follows from $OPT_{g'} \geq OPT_g$. \square

D. FUTURE INDEPENDENT SETTING: CONTINUOUS VERSION

D.1. Utility functions when players have continuous investments

In Section D, we allow investments in resources to be non-discrete. Here, we describe the form of players' utility in the continuous model. Each resource r is associated with a non-increasing value curve $v'_r : \mathbb{R}^{\geq 0} \rightarrow \mathbb{R}^{\geq 0}$. In the future-independent setting,

$$u_i(a_1, \dots, a_n) = \sum_{r=1}^m \int_{x_{i,r}}^{x_{i,r} + a_{i,r}} v'_r(t) dt,$$

where $x_{i,r} = \sum_{i'=1}^{i-1} a_{i',r}$ is the amount already invested in resource r by earlier players.

In this setting, in order to prove a theorem analogous to Theorem 5.5 in the discrete setting, we need an analogue to Lemma 5.1 that holds in the full-information continuous setting. We no longer have the tight connection between our setting and matching; nonetheless, the fact that the greedy strategy 4-approximates to OPT continues to hold

LEMMA D.1. *The greedy strategy for many-to-one online, continuous, resource-weighted "matching", where players arrive online and have tuples of allowable volumes of resources, has a competitive ratio of $\frac{1}{4}$.*




Fig. 1. Suppose the blue regions are those selected by the players who got those regions in OPT , and the red regions are those selected by some other player. Then, if some greedy player(s) have taken at least half of the value of the optimal regions for another player, at least that much utility has been gained by the greedy players. If not, half the value is still available for the player at hand.

PROOF. The proof is identical to the proof of Lemma 5.1, with the exception that we no longer want matchings μ, μ' but rather correspondences between continuous regions of v'_r . See Figure D.1 for a visual proof sketch. \square

With Lemma D.1, the generalization of Theorem D.2 is immediate.

THEOREM D.2. *Suppose that \mathcal{M} is a (α, β, γ) -counter, and that \mathcal{M} is an underestimator. Then, for any continuous, future-independent resource-sharing game g , $CR_{\text{GREEDY}}(\mathcal{M}, g) = O(\alpha\beta)$.*

E. ANALYSIS OF PRIVATE COUNTERS

PROOF OF LEMMA 7.2. We assume the reader is familiar with the TreeSum mechanism. The privacy of this construction follows the same argument as for the original constructions. One can view m independent copies of the TreeSum protocol as a single protocol where the Laplace mechanism is used to release the entire vector of partial sums. Because the ℓ_1 -sensitivity of each partial sum is 1 (since $\|a_t\| \leq 1$), the amount of Laplace noise (per entry) needed to release the m -dimensional vector partial sums case is the same as for a dimensional 1-dimensional counter.

To see why the approximation claims holds, we can apply Lemma 2.8 from [Chan et al. 2011] (a tail bound for sums of independent Laplace random variables) with $b_1 = \dots = b_{\log n} = \log n/\epsilon$, error probability $\delta = \gamma/mn$, $\nu = \frac{(\log n)\sqrt{\log(1/\delta)}}{\epsilon}$ and $\lambda = \frac{(\log n)(\log(1/\delta))}{\epsilon}$, we get that each individual counter estimate $s_t(j)$ has additive error $O(\frac{(\log n)(\log(nm/\gamma))}{\epsilon})$ with probability at least $1 - \gamma/(mn)$. Thus, all $n \cdot m$ estimates satisfy the bound simultaneously with probability at least $1 - \gamma$. \square

PROOF OF LEMMA 7.3. We begin with the proof of privacy. The first phase of the protocol is $\epsilon/2$ -differentially private because it is an instance of the “sparse vector” technique of Hardt and Rothblum [2010] (see also [Roth 2011, Lecture 20] for a self-contained exposition). The second phase of the protocol is $\epsilon/2$ -differentially private by the privacy of TreeSum. Since differential privacy composes, the scheme as a whole is ϵ -differentially private. Note that since we are proving $(\epsilon, 0)$ -differential privacy, it suffices to consider nonadaptive streams; the adaptive privacy definition then follows [Dwork et al. 2010].

We turn to proving the approximation guarantee. Note that the each of the Laplace noise variables added in phase 1 of the algorithm (to compute $x_{t,r}$ and τ_j) uses parameter $2/\epsilon'$. Taking a union bound over the mn possible times that such noise is added, we see that with probability at least $1 - \gamma/2$, each of these random variables has absolute value at most $O(\frac{\log(nm/\gamma)}{\epsilon'})$. Since $\frac{2}{\epsilon'} = O(\frac{mk}{\epsilon})$ and $k = O(\log \log(\frac{nm}{\gamma}) + \log \frac{1}{\epsilon})$, we get that each of these noise variables has absolute value $\tilde{O}_\alpha(\frac{m \log(nm/\gamma)}{\epsilon})$ with probability all but $\gamma/2$. We denote this bound E_1 .

Thus, for each counter, the i -th flag is raised no earlier than when the value of the counter first exceeds $\alpha^i(\log n) - E_1$, and no later than when the counter first exceeds $\alpha^i(\log n) + E_1$. The very first flag might be raised when counter has value 0. In that case, the additive error of the estimate is $\log n$, which is less than E_1 . Hence, the mechanism’s estimates during the first phase provide an $(\alpha, E_1, \gamma/2)$ -approximation (as desired).

The flag that causes the algorithm to enter the second phase is supposed to be raised when the counter takes the value $A := \alpha^k(\log n) \geq \frac{\alpha}{\alpha-1} \cdot C_{\text{tree}} \cdot \frac{\log(nm/\gamma)}{\epsilon}$; in fact, the counter could be as small as $A - E_1$. After that point, the additive error is due to the TreeSum protocol and is at most $B := C_{\text{tree}} \cdot \log(n) \cdot \log(nm/\gamma)/\epsilon$ (with probability at least $1 - \gamma/2$) by Lemma 7.2. The reported value $s_{i,r}$ thus satisfies

$$s_{i,r} \geq x_{i,r} - B = \frac{1}{\alpha}x_{i,r} + \underbrace{(1 - \frac{1}{\alpha})x_{i,r} - B}_{\text{residual error}}.$$

Since $x_{i,r} \geq A - E_1$, the “residual error” in the equation above is at least $(1 - \frac{1}{\alpha})(A - E_1) - B = -(1 - \frac{1}{\alpha})E_1 \geq -E_1$. Thus, the second phase of the algorithm also provides $(\alpha, E_1, \gamma/2)$ -approximation. With probability $1 - \gamma$, both phases jointly provide a (α, E_1, γ) -approximation, as desired. \square

F. OTHER GAMES

F.1. Unrelated Machine scheduling games

Below, we denote by t_k^* the minimum cost of job k among all the machines and by q_k^* the machine that achieves this minimum.

PROOF OF THEOREM 8.3. Consider any player i , and let the *displayed* load profile she sees be L . Using greedy strategy, she will put her job on machine q that minimizes $L_q + t_{kq}$ and this in particular shall be at most $|L|_\infty + t_k^*$. Since the *true* makespan before this player placed her job is at most $\alpha|L|_\infty + \beta$, hence after she places her job, for the displayed load profile L' , $|L'|_\infty \leq \alpha(\alpha|L|_\infty + \beta + t_k^*) + \beta \leq \alpha^2(|L|_\infty + 2\beta + t_k^*)$.

Using the above reasoning for every player in the sequence, we have the displayed load profile L_n at the end of the sequence has the property that $|L_n|_\infty \leq \alpha^{2n}(|L_0|_\infty + 2n\beta + \sum_{k=1}^n t_k^*)$, where L_0 is the load profile shown to the first player. But $|L_0|_\infty$ is at most β , since the true load on all machines is zero at that point.

Since the displayed makespan at the end of the sequence is at most $\alpha^{2n}(\beta + 2n\beta + \sum_{k=1}^n t_k^*)$, hence the true makespan is at most $\alpha^{2n+1}(\beta + 2n\beta + \sum_{k=1}^n t_k^*) + \beta$. Since $OPT \geq \sum_{k=1}^n t_k^*/m$ we have our result. \square

F.2. Cut games

PROOF OF THEOREM 8.5. [Leme et al. 2012] Consider the choice made by player t when it is her turn to play. Let C_t be the number of neighbors of player t that have adopted a color by the time it her turn to play. Notice that the total number of edges in the graph is $\sum_t C_t$. Furthermore, the greedy strategy ensures that player t gets value at least $C_t/2$. Since the number of edges in the graph is an upper bound on the optimal social welfare, hence we have the greedy strategy achieving a competitive ratio of 2. \square

PROOF OF THEOREM 8.6. Let us analyze the play made by player t when it is her turn to play. Let R_t and B_t be the true counts of red and blue neighbors of t at that time, and without loss of generality let $R_t \geq B_t$. Either the player chooses the blue color and this guarantees her utility of $C_t/2$, where $C_t = R_t + B_t$. On the other hand, if the player were to choose the color red, it must be the case that the displayed value of the blue counter is at least the displayed value of the red counter. For this to be true, it must be the case that $\alpha B_t + \beta \geq R_t/\alpha - \beta$, and therefore $B_t \geq R_t/\alpha^2 - 2\beta/\alpha \geq C_t/(2\alpha^2) - 2\beta/\alpha$. Hence, in either case, the player achieves utility of at least $C_t/(2\alpha^2) - 2\beta/\alpha$.

Following the analysis used in the proof of Theorem 8.5, we have the result. \square

PROOF OF THEOREM 8.4. Consider the graph to be a long cycle with $2n$ nodes. For ease of analysis, number the nodes 0 through $2n - 1$ with the node numbered i have its neighbors $(i - 1) \bmod 2n$ and $(i + 1) \bmod 2n$. The optimal social welfare is $4n$ obtained by coloring all even numbered nodes with red and the rest with blue.

Consider the sequence of nodes where nodes arrive in the increasing order 0 through $2n - 1$. We claim that through a series of undominated strategy plays on part of each player, the coloring where node $2n - 1$ is colored red and the rest colored blue is achievable. Note that this coloring gives a social welfare of 4.

We now prove our claim. It is an undominated strategy for node 0 to choose the color blue. Node 1 sees one of its neighbors colored blue and the other uncolored. It is an undominated strategy for node 1 to choose color blue as well. This continues and each node until node $2n - 1$ is colored blue. Node $2n - 1$ has both its neighbors colored blue, and so the only undominated strategy for her is to play red. \square

